

TI 2018-027/III
Tinbergen Institute Discussion Paper



Unobserved Components with Stochastic Volatility in U.S. Inflation: Estimation and Signal Extraction

Mengheng Li¹
Siem Jan (S.J.) Koopman²

1: VU Amsterdam

2: VU Amsterdam; Tinbergen Institute, The Netherlands

Tinbergen Institute is the graduate school and research institute in economics of Erasmus University Rotterdam, the University of Amsterdam and VU University Amsterdam.

Contact: discussionpapers@tinbergen.nl

More TI discussion papers can be downloaded at the [Tinbergen Site](#)

Tinbergen Institute has two locations:

Tinbergen Institute Amsterdam
Gustav Mahlerplein 117
1082 MS Amsterdam
The Netherlands
Tel.: +31(0)20 598 4580

Tinbergen Institute Rotterdam
Burg. Oudlaan 50
3062 PA Rotterdam
The Netherlands
Tel.: +31(0)10 408 8900

Unobserved Components with Stochastic Volatility in U.S. Inflation: Estimation and Signal Extraction

Mengheng Li^{1,3,4} and Siem Jan Koopman^{1,2,3*}

¹ *Department of Econometrics, Vrije Universiteit Amsterdam, The Netherlands*

² *CREATES, Aarhus University, Denmark*

³ *Tinbergen Institute Amsterdam, The Netherlands*

⁴ *Economics and Research Division, De Nederlandsche Bank, The Netherlands*

Abstract

We consider unobserved components time series models where the components are stochastically evolving over time and are subject to stochastic volatility. It enables the disentanglement of dynamic structures in both the mean and the variance of the observed time series. We develop a simulated maximum likelihood estimation method based on importance sampling and assess its performance in a Monte Carlo study. This modelling framework with trend, seasonal and irregular components is applied to quarterly and monthly US inflation in an empirical study. We find that the persistence of quarterly inflation has increased during the 2008 financial crisis while it has recently returned to its pre-crisis level. The extracted volatility pattern for the trend component can be associated with the energy shocks in the 1970s while that for the irregular component responds to the monetary regime changes from the 1980s. The scale of the changes in the seasonal component has been largest during the beginning of the 1990s. We finally present empirical evidence of relative improvements in the accuracies of point and density forecasts for monthly US inflation.

Keywords: *Importance Sampling, Kalman Filter, Monte Carlo Simulation, Stochastic Volatility, Unobserved Components Time Series Model, Inflation*

JEL Classification: C32; C53; E31; E37

*Corresponding author: SJ Koopman, Department of Econometrics, Vrije Universiteit Amsterdam School of Business and Economics, De Boelelaan 1105, 1081 HV Amsterdam, The Netherlands. Email s.j.koopman@vu.nl Telephone +31205986019 Fax +31205986020. A supplementary appendix with more material and results is available at <http://sjkoopman.net/papers/LK2018Appendix.pdf>. SJK would like to acknowledge financial support by the Center for Research in Econometric Analysis of Time Series (DNERF78), CREATES, funded by the Danish National Research Foundation.

1 Introduction

An in-depth understanding of the dynamic features of inflation is crucial for monetary policy making and accurate forecasting of inflation. As discussed in [Clarida, Gali, and Gertler \(2000\)](#), the Federal Reserve Bank determines the target nominal interest rate on a forward-looking Taylor rule which has inflation expectations as key variables. Since the 2008 financial crisis, major advanced economies have experienced a prolonged period of low inflation. Many researches address the question whether inflation expectation has become more reluctant (see, for example [Cúrdia et al. 2015](#), [Constâncio 2016](#), [Lawrence and Rudebusch 2016](#) and [Congressional Budget Office 2016](#)). Additionally, due to the sluggish inflation response to accommodative monetary policy in the past years, much attention has been paid to the estimation of the natural rate of interest and interest rate gap; the former is an important policy benchmark while the latter is a monetary coincidental indicator. Yet a commonly found problem in recent literature is that econometric models are very sensitive to the choice of inflation expectation (see, for example [Laubach and Williams 2003](#), [Holston, Laubach, and Williams 2017](#), [Cecchetti, Hooper, Kashyap, and Schoenholtz 2017](#) and [Del Negro, Giannone, Giannoni, and Tambalotti 2017](#)).

These concerns highlight the need for modelling the salient dynamic features of inflation and for producing reliable forecasts which closely measure inflation expectation. In our study we explore the role of volatility in the modelling and forecasting of inflation in an unobserved component (UC) model framework. The UC models have been extensively used for modelling the time-varying mean and extracting dynamic features in macroeconomic time series, see [Harvey \(1989\)](#). It is widely recognised that time-varying volatility should play an important role in the modelling of inflation because it relates directly to the measurement of core inflation in conducting monetary policy. In departure from the Federal Reserve Bank, various central banks focus on the headline instead of the core inflation¹. The reasons are multifold; but central banks tend to recognise that headline inflation without energy and food price changes cannot serve as a good measure of the underlying core inflation. Those components are no longer the only volatile ones due to expansion of international trade and

¹For example, the operational policy of the European Central Bank and the Bank of England is to adjust the policy rate such that the headline inflation forecast or expectation is moving close to its 2% target without inducing instability in the real economy ([Bank of England, 2013](#)).

increasing international price pass-through (Erceg, Henderson, and Levin 2000 and Federal Open Market Committee 2009). Since different measures of core inflation give different policy implications, a fundamental requirement of modelling and forecasting inflation is thus on the dynamics of its volatility and the ability of detecting volatility shocks.

Our general UC model with stochastic volatility (SV) components is referred to as the UCSV model. It can simultaneously capture stylised facts such as time-varying mean, volatility clustering, higher moment autocorrelation, and heavy tailed densities. UCSV is also adopted in the study of Stock and Watson (2007) where they decompose quarterly US inflation into a random walk process and idiosyncratic noise. The model without SV is known as *the local level model* and is advocated by Harvey (1989). The local level model has a fixed *signal-to-noise ratio* (SNR) and thus implies a constant persistence, or a constant discounting of past observations. By letting both the random walk innovations and the idiosyncratic noise be subject to SV, Stock and Watson (2007) analyse the quarterly US headline inflation and extract the core component based on which forecasts are made. Shephard (2015) revisits this model and provides a full Bayesian estimation method. He emphasises that the two SV processes indicate a time-varying SNR whose path reflects the changing memory and persistence of inflation, responding to energy price shocks and changes in monetary policy. In Stock and Watson (2008) it is shown that UCSV captures the stochastic nature of the mean, volatility and persistence of inflation effectively. It significantly outperforms many other existing models in terms of forecasting accuracy, including the “triangle” model of Gordon (1990) and the random walk model of Atkeson and Ohanian (2001).

Despite the basic model structure, parameter estimation is a challenging exercise. For example, Stock and Watson (2007) set the only scaling parameter in their local level model with SV to 0.2 as it delivers good point forecasts. As is known, SV models are extensively studied in financial econometrics, but models with both stochastic mean and stochastic volatility are more relevant for macroeconomic time series. For a detailed discussion about SV models, see the survey by Shephard (2005) and Platanioti, McCoy, and Stephens (2005). Many simulation-based methods have been developed for estimating SV models, most of which uses Bayesian Markov chain Monte Carlo (MCMC) method. Other methods are simulation-based generalised method of moments, indirect inference and quasi-maximum likelihood; see the overview in Platanioti et al. (2005) and references therein. UCSV models

are only conditionally linear, meaning that the likelihood function cannot be obtained via prediction error decomposition with the presence of SV. Shephard (2015) proposes an MCMC method based on the particle filter. His method of computing the marginal likelihood is closely related to the mixture Kalman filter in Chen and Liu (2000) which is essentially a combination of particle filtering and Kalman recursions. To the best of our knowledge, only Bayesian inference is available for UCSV models. In our study we provide a classical maximum likelihood (ML) estimation procedure based on importance sampling. We build on the earlier contributions by Shephard and Pitt (1997), Durbin and Koopman (1997), Richard and Zhang (2007) and Koopman, Lucas, and Scharth (2015). The first two contributions construct an importance sampler from a local second-order Taylor expansion, whereas the last two contributions develop importance sampling methods which are globally efficient. In particular, the numerically accelerated importance sampling (NAIS) of Koopman et al. (2015) finds an efficient importance density from which samples of the unobserved SV process can be drawn in a computationally efficient manner. To obtain an ML estimation method for UCSV models, we modify and extend the NAIS method to multivariate settings and cases where the *conditionally independence* assumption fails². Our proposed simulated ML estimation method is shown to be efficient and valid in terms that the central limit theorem of Geweke (1989) applies. The method is likelihood-based and standard asymptotic results apply. Hence statistical inference for UCSV models relies on standard results.

The paper is organised as follows. Section 2 formulates the UCSV models. Section 3 provides briefly discusses how to apply the proposed simulated ML method. Section 4 studies a Monte Carlo experiment providing evidence that a central limit theorem is valid for UCSV models estimated using our method. In Section 5, we carry out empirical studies in which the model in Stock and Watson (2007) is revisited and we further extend it to a local level plus seasonal model with SV to study the monthly US inflation. We present some interesting findings and comparisons of forecasting ability among model variants. Section 6 concludes. A supplementary appendix which details the importance sampling estimation procedure and lists additional empirical results with robustness checks is available online³.

²For example in SV models, given the unobserved SV process, the remainder is conditionally independent and the log-likelihood can be evaluated straightforwardly; see Durbin and Koopman (1997).

³We refer to <http://sjkoopman.net/papers/LK2018Appendix.pdf> for supplementary material.

2 Unobserved components with stochastic volatility

2.1 Local level model with stochastic volatility

Before we introduce our modelling framework in its generality, we discuss the unobserved components time series model for US quarterly inflation as considered by [Stock and Watson \(2007, 2008\)](#). The observed time series of inflation is denoted by y_t with time index $t = 1, \dots, n$ where n is the length of the time series. The local level model decomposes the time series variable into two unobserved components: the level μ_t that is modelled as a random walk process, and the irregular noise ϵ_t that is a serially and mutually independent process. The random walk process and the white noise process are both scaled by separate SV processes h_t^μ and h_t^y , respectively. Let $N(a, b)$ denote a normal distribution with mean a and variance b . We obtain the model

$$\begin{aligned} y_t &= \mu_t + \exp(h_t^y / 2) \epsilon_t, & \epsilon_t &\sim N(0, 1), \\ \mu_{t+1} &= \mu_t + \exp(h_t^\mu / 2) \eta_t^\mu, & \eta_t^\mu &\sim N(0, 1), \end{aligned} \tag{1}$$

where the SV components h_t^y and h_t^μ can be interpreted as log-variance components for the two Gaussian disturbances in the model. We will term $\exp(h_t^y/2)$ the transitory volatility and $\exp(h_t^\mu/2)$ the permanent volatility; because the former affects y_t only at time t , while the latter accumulates due to the random walk process. We assume that ϵ_t and η_t^μ are serially and mutually independent for all t . The SV components are modelled as random walks with possibly correlated innovations, we have

$$\begin{aligned} h_{t+1}^y &= h_t^y + \sigma_y \zeta_t^y, & \zeta_t^y &\sim N(0, 1), \\ h_{t+1}^\mu &= h_t^\mu + \sigma_\mu \zeta_t^\mu, & \zeta_t^\mu &\sim N(0, 1), & E(\zeta_t^y \zeta_t^\mu) &= \rho, \end{aligned} \tag{2}$$

where $\sigma_y > 0$ and $\sigma_\mu > 0$ are often referred to as volatility of volatility (VoV) coefficients and $-1 < \rho < 1$ is a correlation coefficient. In this specification, the innovations for the two SV processes are correlated but conditional on the SV processes, the measurement equation for y_t and the updating equation for μ_t are independent. Furthermore, the SV processes in (2) are independent of both the measurement and updating equations in (1).

2.2 Properties of local level model without stochastic volatility

When we assume that SV processes in (2) are not present, the Stock and Watson model (1) reduces to the local level model of Harvey (1989) as given by

$$y_t = \mu_t + \sigma_y \epsilon_t, \quad \mu_{t+1} = \mu_t + \sigma_\mu \eta_t^\mu. \quad (3)$$

Here $\sigma_y > 0$ and $\sigma_\mu > 0$ denote the constant standard deviations of the disturbances in the measurement and updating equations, respectively. We define the *signal-to-noise ratio* (SNR) as $q = \sigma_\mu^2 / \sigma_y^2$. The celebrated Kalman filter computes the prediction of μ_t given the past observations y_1, \dots, y_{t-1} , and given a value of q , that is the minimum mean squared error $\hat{\mu}_{t|t-1}$, together with its mean squared error $\sigma_\epsilon^2 p_{t|t-1}$, via the recursion

$$\hat{\mu}_{t+1|t} = \hat{\mu}_{t|t-1} + \frac{p_{t|t-1}}{p_{t|t-1} + 1} (y_t - \hat{\mu}_{t|t-1}), \quad p_{t+1|t} = p_{t|t-1} - p_{t|t-1} \frac{p_{t|t-1}}{p_{t|t-1} + 1} + q, \quad (4)$$

with $\hat{\mu}_{2|1} = y_1$ and $p_{2|1} = 1$, for $t = 2, \dots, n$; see Harvey (1989) and Durbin and Koopman (2012) for derivations of these equations. The SNR q is central to these equations as it determines the value of $p_{t|t-1}$ and therefore also the value of $\hat{\mu}_{t|t-1}$. It is known that the local level model rationalises the *exponentially weighted moving average* (EWMA) weighting scheme for the one-step ahead forecasting of y_t that is given by $\tilde{y}_t = \tilde{\mu}_{t|t-1}$ where

$$\tilde{\mu}_{t+1|t} = (1 - \lambda) \tilde{\mu}_{t|t-1} + \lambda y_t, \quad (5)$$

with constant $0 < \lambda < 1$. We can set λ equal to the observation weight implied by the *steady state* Kalman filter that is obtained when we let $t \rightarrow \infty$ such that $p_{t|t-1} = p_{t+1|t} = \bar{p}$. Solving the limiting case for the $p_{t+1|t}$ equation in (4), we obtain $\bar{p} = -0.5q + 0.5\sqrt{q^2 + 4q}$. Given this steady state solution, the update $\hat{\mu}_{t+1|t}$ can be represented as (5) with $\lambda = (\bar{p} + q) / (\bar{p} + q + 1) = (q + \sqrt{q^2 + 4q}) / (2 + q + \sqrt{q^2 + 4q})$. Hence the Kalman filter (4) for the local level model (3) converges to the EWMA forecasting scheme.

The forecast function of the local level model (3) is also equivalent to the forecast function of the first-order integrated moving average process, that is the IMA(1) process, but with a

restricted MA coefficient θ . The IMA(1) model can be given by

$$\Delta y_t = \xi_t + \theta \xi_{t-1}, \quad \xi_t \sim N(0, \sigma^2), \quad (6)$$

with difference operator Δ , MA coefficient θ and variance $\sigma^2 > 0$. By equating the first-order autocorrelations of the local level and the IMA(1) models, θ can be expressed as a function of the SNR q , that is $\theta = (\sqrt{q^2 + 4q} - 2 - q) / 2$. It follows that $-1 < \theta < 0$. Further it can be shown that $\theta = \lambda - 1$, when accounting for the invertibility property of the MA coefficient. The mean reversion or persistence of Δy_t solely depends on the SNR q . We can conclude from these properties of the local level model that the SNR q controls both the weights in the EWMA forecast function and the moving average coefficient θ .

The memory index is introduced by Shephard (2015) who refers to it as the *life span* of the time series process. It measures the number of past observations in the associated EWMA forecasting function (5), as represented by the exponentially decaying weighted sum

$$\tilde{\mu}_{t+1|t} = \lambda \sum_{j=0}^{\infty} (1 - \lambda)^j y_{t-j},$$

that have a value that is larger than 0.1. A past observation that is weighted by a value smaller than 0.1 is considered to have not much impact on the forecast \tilde{y}_t , which in our case reflects the inflation expectation. More specifically, the memory index is defined as the value of m such that $(1 - \lambda)^m = 0.1$. It follows immediately that

$$m = \log(0.1) / \log(1 - \lambda) = \log(0.1) / \log(-\theta), \quad 0 < m < \infty. \quad (7)$$

The memory index solely depends (indirectly) on the SNR q . When the memory index m is small (due to a large value of q), the time series has short memory and is not persistent. When m equals, say, 20, the time series has memory and is persistent since many past observations (about 20) have a considerable impact on the forecast. The memory index has a direct and straightforward interpretation.

2.3 Properties of local level model with stochastic volatility

The insights obtained from the properties discussed above are also applicable to the model (1)-(2). The SV components let the SNR q , the EWMA weight λ and the MA coefficient θ become a time-varying process. The additional feature of the model with SV is that it simultaneously allows for a stochastically time-varying mean, stochastic volatility, and time-varying persistence. Hence it is capable of reconciling the conflicting views of [Marques et al. \(2004\)](#), [Cecchetti and Debelle \(2006\)](#), and [Pivetta and Reis \(2007\)](#)⁴. We discuss these features in more detail in Section 5.

For completeness, we define the various time-varying entities implied by the local level model with SV, in terms of the processes h_t^z , with $z = y, \mu$, that is

$$q_t = \exp(h_t^\mu - h_t^y), \quad \lambda_t = \frac{q_t + \sqrt{q_t^2 + 4q_t}}{2 + q_t + \sqrt{q_t^2 + 4q_t}}, \quad \theta_t = \lambda_t - 1, \quad m_t = \frac{\log(0.1)}{\log(1 - \lambda_t)},$$

where q_t is the time-varying SNR, λ_t is the time-varying coefficient of the EWMA forecasting function, θ_t is the time-varying MA coefficient of the IMA(1) representation and m_t is the time-varying memory index, all implied by the model (1)-(2). Since all these entities are time-varying, the associated properties of the time series generated by (1)-(2) are locally defined. The time-varying memory index m_t indicates the extend to which the properties are locally defined.

2.4 Unobserved components with stochastic volatility

Next we consider the extension of the local level model with SV (1)-(2) of including a dynamic seasonal component in the measurement equation, in order to account for seasonal effects in the time series of inflation (when it is not seasonally adjusted). We also explore another departure from the model of [Stock and Watson \(2007\)](#) that has the SV processes modelled as correlated random walks. We consider correlated SV components that are modelled as stationary autoregressive processes with non-zero unconditional means. The full specification of our unobserved components model with stochastic volatility (UCSV) model is given below.

⁴[Cecchetti and Debelle \(2006\)](#) find that the most significant change of inflation process in the last three decades is in the mean; and once this has been accounted for persistence is found to be low. [Pivetta and Reis \(2007\)](#) conclude that inflation persistence has remained high until the financial crisis. [Marques et al. \(2004\)](#) argue that the determination of inflation persistence depends on the chosen model for the mean.

For time series of inflation y_t , we consider the measurement equation $y_t = \mu_t + \gamma_t + \epsilon_t$ where μ_t and γ_t are the level and seasonal components, respectively, with their dynamic specifications given as follows. (a) The *level* or *trend* component μ_t is given by (1) but with the default option that the SV random walk processes are replaced by autoregressive processes as given by

$$h_{t+1}^z = (1 - \phi^z) \alpha^z + \phi^z h_t^z + \sqrt{1 - \phi^{z2}} \sigma_z \zeta_t^z, \quad \zeta_t^z \sim N(0, 1), \quad (8)$$

for $z = y, \mu$, where α^z is a constant, $0 < \phi^z < 1$ is an autoregressive coefficient, and where ζ_t^y and ζ_t^μ can be contemporaneously dependent, with correlation coefficient $\rho^{y\mu}$.

(b) The dynamic *seasonal* component is given by

$$\gamma_t + \gamma_{t-1} + \dots + \gamma_{t-s+1} = \exp(h_t^\gamma / 2) \eta_t^\gamma, \quad \eta_t^\gamma \sim N(0, 1), \quad (9)$$

where s is the seasonal length ($s = 4$ for quarterly data and $s = 12$ for monthly data) and the specification for the SV scaling h_t^γ is given by (8) with $z = \gamma$ and we can possibly have $\rho^{y\gamma} \neq 0$ and $\rho^{\mu\gamma} \neq 0$. The non-zero correlations can be particularly interesting in economic analyses since volatility responses to macro shocks are often correlated. The seasonal component is clearly only relevant in our modeling framework when economic time series are not seasonally adjusted. The specification (9) is referred to as the seasonal dummy specification and it states that the sum of the seasonal effects is zero in expectation. Alternative specifications for the seasonal component can also be considered; see the discussion in [Harvey \(1989\)](#).

The time-varying SNR for the model extended with a seasonal component can also be composed by two SNRs as given by

$$q_t^\mu = \exp(h_t^\mu - h_t^y), \quad q_t^\gamma = \exp(h_t^\gamma - h_t^y),$$

for $t = 1, \dots, n$. The properties of the time series processes rely on the specification of the components and the paths of their corresponding time-varying SNRs. Although for many purposes the properties of the seasonal component can be of particular interest, in our current study most of the focus will be on the level component and its properties that are mostly determined by the SNR q_t^μ and its evolution over time.

Since the SV components rely on linear Gaussian autoregressive processes, it follows that the SNR processes have well-defined first and second moments as implied by the properties of the log-normal distribution and are given by

$$E(q_t^z) = \exp(\alpha^z + \sigma_{yz}^2 / 2), \quad \text{Var}(q_t^z) = [\exp(\sigma_{yz}^2) - 1] \exp(2\alpha^z + \sigma_{yz}^2), \quad (10)$$

where $\sigma_{yz}^2 = \sigma_y^2 + \sigma_z^2 + 2\sigma_y\sigma_z\rho^{yz}$, for $z = \mu, \gamma$. The two moments provide information about the global dynamic properties of the UCSV model as discussed above.

In general, the measurement equation can be extended with more dynamic components. Since many macroeconomic time series are subject to the real economy and its business cycle features, it is of particular interest to consider a stationary autoregressive process with SV as an additional dynamic component (an example is given in the supplementary appendix). Furthermore, we can extend the UCSV model with regression effects that have time-varying coefficients with SV components. Also, the level component can be modified to include a stochastic slope or gradient specification with an SV component. In the supplementary appendix, we discuss such generalisations within the general state space framework of [Durbin and Koopman \(2012\)](#). The consequences for our proposed simulated maximum likelihood estimation procedure are also discussed. Finally, we show in the supplementary appendix that our UCSV model can be regarded as a special case of existing models in the econometric and macroeconomic literatures; the extension with SV components and the corresponding estimation methodology is therefore relevant from a much wider perspective.

3 Simulated maximum likelihood estimation

3.1 Simulated likelihood function

Importance sampling is often used to evaluate intractable densities but it has also been used for the evaluation of the likelihood function for SV models; see [Danielsson \(1994\)](#), [Shephard and Pitt \(1997\)](#), [Durbin and Koopman \(1997\)](#), and [Durham and Gallant \(2002\)](#). Here we adopt the main principles of this likelihood-based approach for the model as described in [Section 2](#) and by a modification of the numerically accelerated importance sampling (NAIS) method of [Koopman et al. \(2015\)](#) with the purpose to treat multiple SV components in

the model, not only one. We collect the parameters of the UCSV model in the vector ψ ; for the inflation model, it consists of the coefficients α^z , ϕ^z and σ_z , for $z = y, \mu, \gamma$. Our main concern is the evaluation and the maximisation (with respect to ψ) of the likelihood function. We provide the details and report various numerical and statistical improvements of our approach. The method essentially relies on using regression computations to adopt an importance density and averaging over simulations from the importance density.

The density $p(\cdot)$ refers to the model and $g(\cdot)$ represents the importance density. We typically take a linear Gaussian state space model as an approximation of the true model and which is also used as the importance density $g(\cdot)$. We define $k \times 1$ vector h_t as the collection of the k SV components in the model at time t ; for example, for model (1)-(2) we have $k = 2$ and vector $h_t = (h_t^y, h_t^\mu)'$. The $k \times n$ matrix $h_{1:n} = (h_1, \dots, h_n)$ contains all values of the SV components in the model. The simulated log-likelihood function $\log L(\psi)$ for the realised observations $y_{1:n} = (y_1, \dots, y_n)$ can be constructed via importance sampling and, after a standard bias-correction, is given by

$$\log L(\psi) = \log L_g(\psi) + \log \bar{\omega} + s_\omega^2 / (2M\bar{\omega}^2),$$

where $\log L_g(\psi)$ is the Gaussian log-likelihood function that is computed via the prediction error decomposition using the Kalman filter applied to the approximating linear Gaussian state space model, $\bar{\omega}$ and s_ω^2 are the mean and variance, respectively, of the importance weight function as given by

$$\omega(h; \psi) = p(y_{1:n}|h_{1:n}; \psi) / g(y_{1:n}|h_{1:n}; \psi),$$

and M is the number of simulations taken from the importance density, that is $h^i \sim g(h_{1:n}|y_{1:n}; \psi)$ for $i = 1, \dots, M$; see the development and discussions in [Durbin and Koopman \(1997\)](#). The simulated estimate of the log-likelihood function is then computed as

$$\log \widehat{L}(\psi) = \log L_g(\psi) + \log \widehat{\omega} + \widehat{s}_\omega^2 / (2M\widehat{\omega}^2), \quad (11)$$

where $\widehat{\omega} = M^{-1} \sum_{i=1}^M \omega^{(i)}$, $\widehat{s}_\omega^2 = (M-1)^{-1} \sum_{i=1}^M (\omega^{(i)} - \widehat{\omega})^2$, $\omega^{(i)} = \omega(h^i; \psi)$, for $i = 1, \dots, M$.

When the same random numbers are used for the sampling of h^i and the computation of

the loglikelihood estimate (11) for different values of ψ , we can maximise (11) with respect to the parameter vector ψ . This is the key aspect of parameter estimation for the local level model with SV; see also the discussions in [Stock and Watson \(2007\)](#) and [Shephard \(2015\)](#). Our approach of using importance sampling for estimation and signal extraction is novel and their details are given below. In the Appendix, we provide the details of the importance sampling estimation of latent processes in the UCSV model including those for the SV components. A particle filtering algorithm is also implemented for the UCSV model for real-time filtering and for computing one-step ahead prediction errors and associated diagnostic tests, see the supplementary appendix for details.

3.2 The choice of importance density: the NAIS method

Building on the earlier work on importance sampling, and in particular to the efficient importance sampling (EIS) method of [Richard and Zhang \(2007\)](#), the numerically accelerated importance sampling (NAIS) method, as developed by [Koopman et al. \(2015\)](#), is shown to be effective in obtaining an importance density for models such as (1)-(2). The EIS and NAIS methods are based on finding an importance density that we can represent as

$$g(y_t|h_t; \psi) = \exp\left(r_t + b_t'h_t - \frac{1}{2}h_t'C_t h_t\right), \quad (12)$$

where r_t is an integrating constant while $k \times 1$ vector b_t and $k \times k$ matrix C_t , for $t = 1, 2, \dots, n$, are implicitly defined as functions of data $y_{1:n}$ and parameter vector ψ . Given that r_t is chosen such that $g(y_t|h_t; \psi)$ integrates to one, the construction of an importance density in this framework is reduced to the choice of a set of importance sampling parameters $b_{1:n}, C_{1:n}$.

It is shown by [Koopman et al. \(2015\)](#) for $k = 1$ that (12) is equivalent to the conditional density $g(y_t^+|h_t; \psi)$ with $y_t^+ = C_t^{-1}b_t$ that is modelled by a linear Gaussian state space model with the observation equation

$$y_t^+ = h_t + \epsilon_t^+, \quad \epsilon_t^+ \sim N(0, C_t^{-1}), \quad t = 1, 2, \dots, n, \quad (13)$$

and with the state equation for h_t represented by the autoregressive process as formulated in (8). A set of optimal importance sampling parameters $b_{1:n}$ and $C_{1:n}$ can be taken as the

solutions of the variance minimisation problem given by

$$\min_{b_t, C_t} \int \lambda_t^2(y_t, h_t; \psi) \omega_t(y_t, h_t; \psi) g(h_t | y_{1:n}^+; \psi) dh_t, \quad \omega_t(y_t, h_t; \psi) = \frac{p(y_t | h_t; \psi)}{g(y_t^+ | h_t; \psi)}, \quad (14)$$

where $\lambda_t(y_t, h_t; \psi) = \log \omega_t(y_t, h_t; \psi) = \log p(y_t | h_t; \psi) - r_t - b_t' h_t + \frac{1}{2} h_t' C_t h_t$, for $t = 1, \dots, n$. For each t , the evaluation of the integral, and its subsequent analysis, can be based on simulation (EIS) or on numerical integration (NAIS). The latter method has our preference because it is numerically fast and robust. Since the integral in (14) can be numerically approximated with high precision as a Gauss-Hermite computation, the basic NAIS method constitutes, for each t , the minimisation as given by

$$\min_{b_t, C_t} \sum_{j=1}^K w_{tj} \lambda_t^2(y_t, \tilde{h}_{tj}; \psi), \quad w_{tj} = w^*(z_j) \exp(z_j^2 / 2) \omega_t(y_t, \tilde{h}_{tj}; \psi) / \sqrt{2\pi},$$

with integer $K > 1$, abscissae or node z_j and associated Gauss-Hermite weight $w^*(z_j)$, for $j = 1, \dots, K$. For a given value of K , we typically have $K = 12$, both z_j and $w^*(z_j)$ are tabulated. We further have

$$\tilde{h}_{tj} = \hat{h}_t + \sqrt{V_t} z_j, \quad \hat{h}_t = E_g(h_t | y_{1:n}^+; \psi), \quad V_t = E_g(h_t - \hat{h}_t | y_{1:n}^+; \psi)(h_t - \hat{h}_t | y_{1:n}^+; \psi)',$$

where E_g refers to expectation with respect to the importance model density $g(y_t^+ | h_t; \psi)$, that is the approximating linear Gaussian state space model with observation equation (13). The smoothed mean \hat{h}_t and corresponding variance V_t are computed via the Kalman filter and an associated smoothing algorithm applied to the approximating model; for example, see Durbin and Koopman (2012, Chapter 4). The solution to this minimisation for each t is a weighted regression with K ‘‘observations’’, dependent variable $\log p(y_t | h_{tj}; \psi)$, explanatory variables \tilde{h}_{tj} and \tilde{h}_{tj}^2 , weight w_j . The least squares estimates for the coefficients corresponding to \tilde{h}_{tj} and \tilde{h}_{tj}^2 are the solutions for b_t and C_t , respectively. This regression can only be carried out with existing values for $b_{1:n}$ and $C_{1:n}$ as they are needed for the computation of \tilde{h}_{tj} via the approximating model (13). The overall solution of (14) is therefore an iterative process: for given values of $b_{1:n}$ and $C_{1:n}$, we obtain new values for these variables via the weighted regression. After a small number of iterations, convergence to a solution is achieved.

3.3 The importance density for the UCSV model : modified NAIS

The implementation of NAIS for the UCSV model is faced with two imminent challenges. First, the importance sampling method relies on the *conditional independence* assumption which leads to (14). In particular, the implied factorization

$$p(y_{1:n}|h_{1:n}; \psi) = \prod_{t=1}^n p(y_t|h_t; \psi),$$

is not valid because the components (trend and seasonal) in the UCSV model imply a dynamic structure for $y_{1:n}$, even when we condition on the SV components in $h_{1:n}$. Second, we have treated the scalar case of h_t , that is $k = 1$, but for the UCSV model h_t represents a $k \times 1$ vector of SV components; in our case we have $k = 3$ and h_t contains SV elements for the irregular, trend and seasonal components. Hence the integral in (14) refers to a three-dimensional integration and it is implied that

$$\int x(h_t)dh_t = \int \int \int x(h_{1t}, h_{2t}, h_{3t})dh_{1t}dh_{2t}dh_{3t},$$

for some function $x(\cdot)$. The Gauss-Hermite calculations therefore concern a three-dimensional summation which is computationally more demanding; for example, the summation of $K = 12$ becomes a summation of $K^3 = 1,728$.

We address these two challenges via a modification of the NAIS method as follows. The conditional independence is not applicable to $p(y_{1:n}|h_{1:n}; \psi)$ but this density represents the standard (linear Gaussian) UC model with time-varying variances for the disturbances associated with the irregular, trend and seasonal components and given by $h_{1:n}$. The prediction error decomposition can therefore be applied conditional on $h_{1:n}$ and is given by

$$p(y_{1:n}|h_{1:n}; \psi) = p(y_1|h_1; \psi) \prod_{t=2}^n p(y_t|y_{1:t-1}, h_{1:n}; \psi) = \prod_{t=1}^n p(v_t|h_{1:n}; \psi) = \prod_{t=1}^n p(v_t|h_t; \psi),$$

where v_t is the prediction error $v_t = y_t - E(y_t|y_{1:t-1}h_{1:n}; \psi)$ which is normally and serially independent distributed with mean zero and variance F_t ; both v_t and F_t are delivered by the Kalman filter. An example of the Kalman filter is for the local level model as given by (4) with $v_t = y_t - \hat{\mu}_{t|t-1}$ and $F_t = \sigma_\epsilon^2 (p_{t|t-1} + 1)$. The conditional independence assumption

applies to $p(v_{1:n}|h_{1:n}; \psi)$ and in (14) we can replace the importance weight function by

$$\omega_t(y_t, h_t; \psi) = \frac{p(v_t|h_t; \psi)}{g(y_t^+|h_t; \psi)}. \quad (15)$$

Finally we notice that Kalman filter initialisation issues are a concern but that they can be treated in a standard way as discussed in Durbin and Koopman (2012, Chapter 5).

To ease the computational burden of the Gauss-Hermite computations of dimension K^k we follow a pruning method that we illustrate for the two-dimensional case of $k = 2$ for our local level model with SV. The top left panel of Figure 1 presents the two-dimensional Gauss-Hermite quadrature results in a rectangular set of function evaluations at the nodes (z_{j_1}, z_{j_2}) for $j_1, j_2 = 1, \dots, K$. However, many of these node combinations lead to near-zero likelihood values as they are located too far in the tails of the bivariate normal distribution. Hence we adopt a pruning technique that enables us to reduce the number of nodes for which an evaluation is needed. We have set a threshold θ_K which in a two-dimensional case can be given by

$$\theta_K := \frac{w^*(z_{j_1}) \cdot w^*(z_{\lfloor \frac{K+1}{2} \rfloor})}{K},$$

where $\lfloor a \rfloor$ is the floor operator for any scalar a . If we drop all grid points in the Gauss-Hermite quadrature that obtains a weight below θ_K , we obtain the approximation

$$\int x(h_t) dh_t \approx \sum_{j_1=1}^K \sum_{j_2=1}^K \mathbf{1}_{\{w^*(z_{j_1})w^*(z_{j_2}) \geq \theta_K\}} w_{tj_1} w_{tj_2} x(\tilde{h}_{tj_1}, \tilde{h}_{tj_2}),$$

where $x(h_t)$ may represent here the function $\lambda_t^2(y_t, h_t; \psi)$ in (14). The effect of pruning is presented in the top right panel of Figure 1.

Although $h_{1:n}$ may consist of independent rows of SV paths, we take the estimate of $h_{1:n}$ as its mean conditional on $y_{1:n}^+$ with respect to the importance density $g(h_t|y_{1:n}^+; \psi)$ and hence the rows of a $\hat{h}_{1:n}$ are dependent (or correlated). In case of correlation among the function inputs, the rectangular combination of Gauss-Hermite nodes becomes a skewed hyper-parallelogram grid. Jäckel (2005) has shown that in such cases the quadrature comes from a rotated plane as is illustrated in Figure 1; the pruning method remains valid after such rotations.




Figure 1
Two-dimensional GH quadrature

Top left: Combination of GH nodes leads to rectangular set of function evaluations.

Top right: Effect of pruning. Bottom left: Set of function evaluations due to correlation.

Bottom right: Effect of pruning when there is correlation.

4 A Monte Carlo study

Although the likelihood function needs to be estimated, our proposed procedure for likelihood evaluation is exact as it is only subject to simulation error. Next we verify in a Monte Carlo study whether standard asymptotic properties of the simulated maximum likelihood estimator applies in the context of our UCSV modelling framework.

4.1 Model design

In [Stock and Watson \(2007, 2008\)](#), the local level model with SV is adopted for the modelling of US inflation. Their analyses have shown superior forecasting performance of the local level model with SV over many other competitive models. We carry out a similar study for our UCSV model specification in [Section 2.4](#) with only a level component but with SV components for both for the irregular noise ϵ_t and the level disturbances η_t^μ as specified in [\(8\)](#) for $z = y, \mu$, respectively. The data generation process $y_t = \mu_t + \epsilon_t$ requires values for the coefficients in [\(8\)](#), α^z , ϕ^z and σ_z , for $z = y, \mu$, and they are set as follows. The SV

process h_t^y is given a higher persistence than the SV process h_t^μ ; specifically we have set $\phi^y = 0.9$ and $\phi^\mu = 0.6$. To investigate the settings of a high, medium and low SNR, we have set the coefficients α^μ , σ_y , and σ_μ , relative to $\alpha^y = 0$, such that $E(q_t^\mu) = \bar{q}$ and coefficient of variation $\sqrt{\text{Var}(q_t^\mu)} / E(q_t^\mu)$, is fixed at $\sqrt{2}/2$, for the three cases $\bar{q} = 4, 1, 0.25$. The expressions for $E(q_t^\mu)$ and $\text{Var}(q_t^\mu)$ are given in equation (10). This particular design avoids possible distortions due to the fact that $\text{Var}(q_t^\mu)$ increases with $E(q_t^\mu)$. For each of the three cases, 1,000 time series are generated with sample sizes $n = 100$, $n = 500$, and $n = 1,000$. In the estimation procedure, we rely on 10 GH nodes with 200 importance samples drawn to calculate the Monte Carlo estimate of the log-likelihood function. We next discuss our simulation results as reported in Table 1. We have also considered the data generation process $y_t = \mu_t + \epsilon_t$ without the SV components (fixed variances) while the estimation is based on the UCSV model. These results are presented in the supplementary appendix.

4.2 Estimation results

The estimation of the parameter vector $(\alpha^\mu, \phi^\mu, \phi^y, \sigma_\mu, \sigma_y)'$ is based on the NAIS procedure and the maximum likelihood method. To ensure stationarity of the SV processes and strictly positive variances, we consider the transformed parameter vector

$$\theta = \begin{pmatrix} \theta_\alpha^\mu = \alpha^\mu \\ \theta_\phi^\mu = \log(\phi^\mu) - \log(1 - \phi^\mu) \\ \theta_\phi^y = \log(\phi^y) - \log(1 - \phi^y) \\ \theta_\sigma^\mu = \log(\sigma_\mu) \\ \theta_\sigma^y = \log(\sigma_y) \end{pmatrix}, \quad (16)$$

such that $\phi^z = \exp(\theta_\phi^z) / [1 + \exp(\theta_\phi^z)]$ and $\sigma_z = \exp(\theta_\sigma^z)$, for $z = \mu, y$. In Table 1 we report the estimation result for the three different SNR values, $\bar{q} = 4, 1, 0.25$. The reported estimates are sample averages with sample standard deviations for the 1,000 simulated time series. Furthermore, we report averages of three diagnostic statistics and the fraction of rejections at 5% significance level. The diagnostic tests are based on the standardised one-step ahead prediction errors as obtained from the particle filter, see the supplementary appendix for details. The diagnostic tests are the Jarque-Bera normality test, the Box-Ljung Q test for

Table 1
Results Monte Carlo Study UCSV Model

Parameter	True θ	Est. θ , $n = 100$	Est. θ , $n = 500$	Est. θ , $n = 1000$	$n \rightarrow \infty$
$\bar{q} = 4$ – high SNR					
$\alpha^\mu = 1.320$	$\theta_\alpha^\mu = 1.320$	2.467 (2.193)	1.556 (0.973)	1.187 (0.394)	(0.411) [0.0016]
$\phi^\mu = 0.600$	$\theta_\phi^\mu = 0.406$	0.249 (0.488)	0.506 (0.197)	0.356 (0.081)	(0.111) [0.0012]
$\phi^y = 0.900$	$\theta_\phi^y = 2.197$	0.752 (1.254)	1.460 (0.886)	1.811 (0.679)	(0.689) [0.0430]
$\sigma_\mu = 0.210$	$\theta_\sigma^\mu = -1.561$	-0.864 (0.906)	-1.121 (0.629)	-1.608 (0.149)	(0.138) [0.0010]
$\sigma_y = 0.110$	$\theta_\sigma^y = -2.073$	-3.183 (2.864)	-2.478 (1.564)	-1.819 (0.593)	(0.684) [0.0104]
	Normality	4.216 [0.07]	5.871 [0.17]	5.416 [0.16]	
	Box-Ljung	9.331 [0.00]	20.084 [0.00]	32.943 [0.08]	
	H($n/3$)	1.261 [0.33]	0.915 [0.21]	0.801 [0.09]	
	LL	-4.152	-27.861	-61.288	
$\bar{q} = 1$ – medium SNR					
$\alpha^\mu = -0.087$	$\theta_\alpha^\mu = -0.087$	-0.105 (0.174)	-0.120 (0.094)	-0.079 (0.014)	(0.013) [0.0008]
$\phi^\mu = 0.600$	$\theta_\phi^\mu = 0.406$	0.934 (0.603)	0.350 (0.108)	0.449 (0.064)	(0.065) [0.0007]
$\phi^y = 0.900$	$\theta_\phi^y = 2.197$	1.818 (0.781)	2.459 (0.461)	2.158 (0.312)	(0.311) [0.0101]
$\sigma_\mu = 0.220$	$\theta_\sigma^\mu = -1.514$	-1.853 (0.945)	-1.347 (0.172)	-1.509 (0.057)	(0.061) [0.0009]
$\sigma_y = 0.120$	$\theta_\sigma^y = -2.120$	-2.315 (1.235)	-2.084 (0.122)	-2.173 (0.094)	(0.083) [0.0061]
	Normality	3.946 [0.07]	4.413 [0.14]	4.035 [0.14]	
	Box-Ljung	8.840 [0.00]	16.546 [0.00]	24.067 [0.02]	
	H($n/3$)	1.105 [0.27]	0.861 [0.10]	0.734 [0.06]	
	LL	-3.556	-24.183	-52.343	
$\bar{q} = 0.25$ – low SNR					
$\alpha^\mu = -2.100$	$\theta_\alpha^\mu = -2.100$	-3.946 (2.107)	-2.630 (1.575)	-1.957 (0.488)	(0.520) [0.0094]
$\phi^\mu = 0.600$	$\theta_\phi^\mu = 0.406$	-0.334 (1.274)	0.012 (0.009)	0.136 (0.018)	(0.016) [0.0027]
$\phi^y = 0.900$	$\theta_\phi^y = 2.197$	2.812 (0.905)	2.695 (0.710)	2.033 (0.304)	(0.307) [0.0021]
$\sigma_\mu = 0.700$	$\theta_\sigma^\mu = -0.357$	-0.978 (1.374)	-0.706 (0.735)	-0.472 (0.119)	(0.154) [0.0047]
$\sigma_y = 0.380$	$\theta_\sigma^y = -0.968$	-1.249 (0.917)	-0.615 (0.460)	-0.914 (0.086)	(0.081) [0.0096]
	Normality	4.660 [0.09]	6.127 [0.15]	6.230 [0.21]	
	Box-Ljung	9.207 [0.01]	22.379 [0.03]	35.068 [0.16]	
	H($n/3$)	1.864 [0.36]	1.253 [0.24]	0.895 [0.09]	
	LL	-4.892	-31.167	-73.641	

We report true parameter values, sample average of their estimates from 1,000 series, and (in parentheses) sample standard deviations, for three sample sizes. We further report the sample averages of normality, Box-Ljung, heteroskedasticity (H) test statistics and [in brackets] their fractions of rejections based on a 5% significance level. LL is the sample average of maximised log-likelihood values. In the last column, we report (in parentheses) the sample averages of the asymptotic standard errors, based on $n = 1000$, and [in brackets] the Monte Carlo numerical standard deviations, based on 30 different random seeds.

the autocorrelation and the heteroskedasticity F test.

The results in Table 1 indicate that our estimation procedure is overall successful in parameter estimation. The results vary across sample size n and SNR \bar{q} . For example, parameter estimates have larger sample standard deviations for a small sample $n = 100$ in relation to the larger samples $n = 500$ and $n = 1,000$. In case of the high SNR $\bar{q} = 4$, σ_μ for the SV process h_t^μ is larger than its counterpart σ_y . Hence the signal μ_t tends to dominate the observation noise ϵ_t . For such cases, and for small sample size $n = 100$, the estimation procedure fails to deliver correct estimates for both θ_ϕ^y and θ_σ^y . However, for increased sample sizes, the estimates become more accurate. For example, for $n = 500$ and in comparison with the parameters θ_ϕ^y and θ_σ^y for the SV of the observation noise, their counterparts θ_ϕ^μ and θ_σ^μ for the signal are estimated more accurately in terms of their standard deviations. This accuracy increases with n as for $n = 1,000$ all parameter estimates are close to their true ones although θ_α^μ , θ_ϕ^μ and θ_σ^μ are estimated more accurately than θ_ϕ^y and θ_σ^y . The opposite is found for the low SNR $\bar{q} = 0.25$. In this case, observation noise eclipses the signal μ_t , which explains why the parameters related to h_t^y are estimated more accurately. However, when the information about system dynamics is somewhat balanced, with global SNR $\bar{q} = 1$, the accuracies of the estimates are comparable for all sample sizes. For this case, we clearly observe that with an increasing sample size, the standard deviations around the true values become smaller. We finally notice that the value of \bar{q} has also consequences for the computational efficiency of likelihood maximisation. In our experiment, when $\bar{q} = 4$ or $\bar{q} = 0.25$, a larger number of likelihood evaluations is required for maximisation (slow convergence) when compared to the case of $\bar{q} = 1$.

Standard deviations reported in Table 1 are computed from the sample of estimates across replications; they reveal the finite-sample properties of our estimation procedure. To investigate how these properties are related to the standard asymptotic properties of our simulated maximum likelihood estimation procedure, we report the sample averages of the asymptotic standard errors, based on our estimates from sample size $n = 1,000$, and for the three different SNR values. We further present the Monte Carlo numerical standard deviations that we computed from 30 rounds of re-estimation with a different random seed; see, for example, [Richard and Zhang \(2007\)](#). We learn that the standard deviations of the estimates for all considered cases are close to the sample average of the asymptotic

standard errors. In cases of low and high SNR values, the estimates produce slightly larger sample standard deviations when compared to the average asymptotic standard errors. These findings substantiate the use of asymptotic standard errors for parameter estimation in UCSV models using our methods. Another finding is that the numerical standard deviations are small when $\bar{q} = 1$, but relatively large when compared to cases of low or high SNR values. It implies that the numerical estimation procedure is more stable under the more “balanced” SNR value of $\bar{q} = 1$.

The Monte Carlo results for $n = 1,000$ are visualised in Figure 2 with histograms and estimated densities of the ML estimates of the (original, non-transformed) parameters. Since estimation is for a transformed parameter vector, the densities of the estimates are typically asymmetric because of the nonlinear transformations, see equation (16). Each panel in Figure 2 presents a histogram that summarises the 1000 ML estimates for an original parameter (row-wise) and a SNR value that is used for generating the data (column-wise). The finite sample distributions of the parameter estimates for data with high and low SNR values have larger variances and heavier tails than for the case of a “balanced” SNR. Interestingly, parameter ϕ^y for the observation error SV h_t^y appears to be estimated less accurately than ϕ^μ . This difference in precision is probably due to the different levels of persistence in the unobserved SV processes as it leads to different finite sample performances.

The one-step ahead prediction errors for the observed time series y_t are obtained from a standard implementation of the particle filter. In this Monte Carlo study we work under correct model specification and hence the diagnostic test statistics can be used in a standard fashion. In Table 1 we report the sample averages of the Jarque-Bera normality test, the Box-Ljung serial correlation Q test, and the heteroskedasticity test. We find that the normality test tends to be rejected more often when the sample size increases and for either small or large SNR values. Similar conclusions can be made for the Box-Ljung and heteroskedasticity test statistics⁵. In the case of $n = 1,000$, the Box-Ljung test rejects 14% of the sample, in case of the small SNR value, but only leads to 8% rejections in the case of a large SNR value. This finding may imply that the power of the Box-Ljung Q test for UCSV models deteriorates somewhat when the SNR value decreases. The heteroskedasticity test shows

⁵The heteroskedasticity test is denoted as $H(n/3)$ which is a F -test with the null of equal variances between the first one-third and the last one-third of the one-step ahead prediction errors.

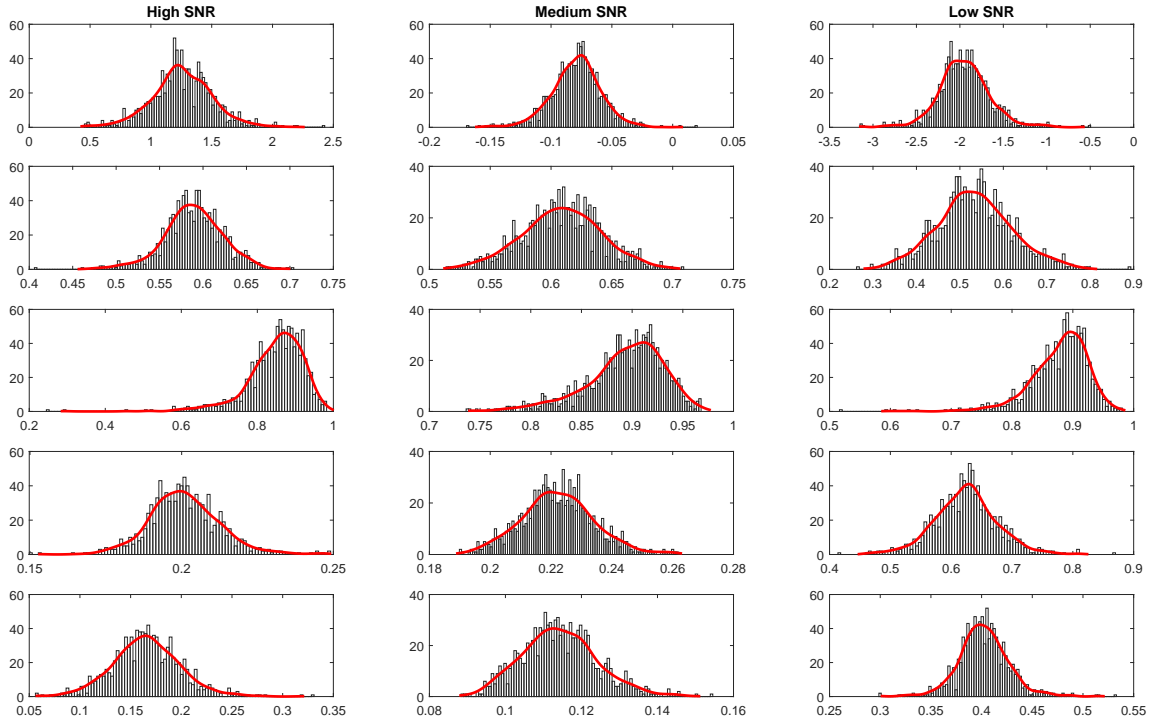


Figure 2

Finite sample distribution of ML estimates for $n = 1,000$.

Column from left to right: high SNR $q = 4$, medium SNR $q = 1$, and low SNR $q = 0.25$.

Row from top to bottom: estimated density of α^μ , ϕ^μ , ϕ^y , σ_μ and σ_y .

overall good and reliable performance, especially for the large sample size of $n = 1000$, and irrespective of the SNR value.

We have also carried out the same Monte Carlo study for our simulated ML procedure but with data generated from the local local model with fixed variances (*without* the SV components). In the supplementary appendix we report the findings in detail. The main finding is that the sample standard deviations of the 1000 ML estimates are large for all parameters, all three sample sizes and all three SNR values. It suggests that most estimates are not significant and hence unreliable. The diagnostic test statistics do not suggest model misspecification and are comparable to those reported in Table 1. We therefore may conclude that our estimation method for the UCSV model is applicable to the case without SV. However, we should report that in case of the small sample size of $n = 100$, there are 92 out of 1000 generated series failing to converge during likelihood maximisation. For the cases of a larger sample size n , the maximisation procedure works satisfactory for all 1000 series. To this extent, our simulated ML method appears to be robust to model misspecification.

Finally, the validity of the simulated maximum likelihood method can be verified on the basis of diagnostic test statistics that rely on extreme value theory and are applied to the importance weights from (15); see [Koopman, Shephard, and Creal \(2009\)](#). In the Appendix we report a selection of these test statistics and it provides sufficient evidence that our simulation method is valid for the encountered SNR values in our study.

5 Modelling and forecasting US inflation

In this section, we first revisit the quarterly headline US inflation analyses of [Stock and Watson \(2007, 2008\)](#) and [Shephard \(2015\)](#), and we compare their results based on MCMC methods with those from our simulated maximum likelihood methods. Next, we model and analyse the monthly US core inflation rate (not seasonally-adjusted), by a UCSV model that includes a stochastic seasonal component with SV; we show that the SVs associated with the different components respond to different economic shocks. We conclude with an extended forecasting study for a range of model specifications; it includes point and density forecasts for different horizons.

5.1 Quarterly headline inflation

[Stock and Watson \(2007, 2008\)](#) have addressed the increasing difficulty of many inflation forecasting models to deliver reliable forecasts, despite the fact that many of these models are based on well-founded economic theory. It is shown that some models outperform others prior to the Volcker-Greenspan scheme in the 1980s, while other models are superior after this period. This finding raises three concerns. Firstly, forecasters need a model which is robust to breaks caused by changes in monetary policy but also in macro-fundamentals (which can be due to energy crisis, financial shocks, etc.). Secondly, a competitive benchmark model is required to make effective comparisons between models and combinations thereof. A policy-maker wants to be assured that the empirical performances of a model are compared with a good-quality benchmark model. Thirdly, inflation forecasting models should be interpretable and consistent throughout a long time period. To reconcile these concerns for the forecasting of the quarterly core inflation, [Stock and Watson \(2007, 2008\)](#) introduce the local level model with SV processes for both disturbances in the model, see equations (1)-(2). It is the basic

UCSV model that we have considered in our Monte Carlo study but with the correlation coefficient set to zero, $\rho = 0$, and the scaling coefficients calibrated at $\sigma_y = \sigma_\mu = 0.2$. In our analysis, we adopt the same UCSV model but estimate the coefficients σ_y , σ_μ and ρ ; these coefficients determine the amount and closeness of the variation of the irregular and signal. For the model without SV we simply report the scaling coefficients for the irregular (σ_y) and the signal (σ_μ). The estimation results are presented in Table 2. The one-step ahead prediction residuals can be computed using the Kalman filter for the UC model without SV and the particle filter for the UCSV model. On the basis of the standardised prediction errors, those for the UCSV model clearly pass the three standard diagnostic tests, for normality, serial correlation and heteroskedasticity, while those for the UC model fail all three tests. The coefficient ρ in the UCSV model is estimated as 0.3 with standard error 0.1; it indicates a moderate yet significant correlation between the SV series h_t^y and h_t^μ . This finding is in accordance to the Bayesian estimates found by Shephard (2015) which we could also reproduce with our implementation of his particle MCMC estimation procedure. However, our ML estimation procedure takes less than 4 seconds while the particle MCMC procedure of Shephard (2015) takes minutes to have a reasonable convergence.

Table 2
UCSV Estimation Results for Quarterly Inflation

	σ_y	σ_μ	ρ	Normality	Box-Ljung	H(n/3)	LL
UCSV	0.349 [0.087]	0.147 [0.064]	0.322 [0.104]	0.615	0.361	0.817	-146.561
UC	0.244 [0.031]	0.373 [0.026]		0.002	0.000	0.021	-201.353

Estimation results for the local level model with SV components (UCSV) and the local level model (UC) without SV. The scaling coefficients σ_y and σ_μ are for the disturbances of the two volatility processes in the UCSV model and represent the standard deviations in the observation and state disturbances, respectively. Standard errors are in square brackets. For the three tests, we present their p -values. LL is the maximised log-likelihood value.

Figure 3 summarises the main fit of the quarterly US inflation series. All graphs are obtained with ψ equal to its ML estimate $\hat{\psi}$. The graph panels (iii) and (iv) display the importance sampling estimates of the time-varying volatility estimates; these are the transitory volatility $E\left(\exp\left(\frac{1}{2}h_t^y\right)|y_{1:n}; \psi = \hat{\psi}\right)$ and the permanent volatility $E\left(\exp\left(\frac{1}{2}h_t^\mu\right)|y_{1:n}; \psi = \hat{\psi}\right)$. It is from the late 1970s, at the time when the oil crises start to take place, that the permanent volatility signal increases significantly. The behaviour of permanent volatility highlights the

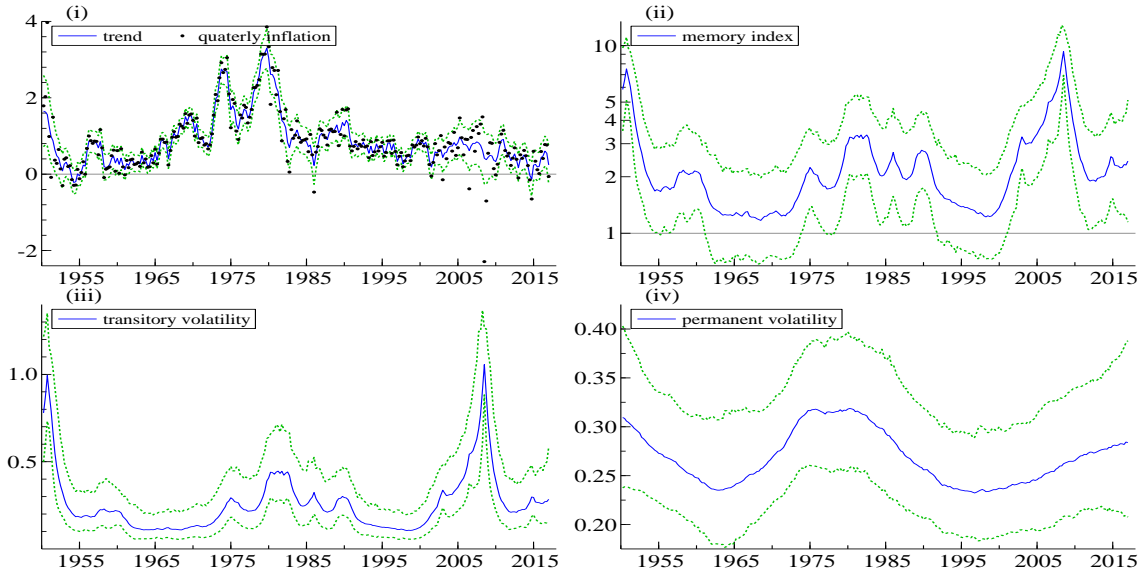


Figure 3

UCSV Decomposition of Quarterly Inflation.

- (i) Inflation series and its trend component μ_t ; (ii) Memory index m_t , defined in Section 2.3;
 (iii) Transitory volatility $\exp(h_t^y/2)$; (iv) Permanent volatility $\exp(h_t^\mu/2)$.

Green dashed lines indicate the 95% confidence bands.

volatility response to energy shocks. Together with a prolonged period of relatively tranquil transitory volatility, with only some increases at the end of the Volcker-Greenspan monetary regime in the early 1980s, we observe a memory index, as displayed in graph panel (ii), that climbs from fewer than one quarter towards three quarters. This finding is supported by Goodfriend and King (1997) and King and Wolman (1999) who argue that the volatility response to monetary shocks leads to reluctant inflation expectation. This becomes even more evident in the 2008 financial crisis. Transitory volatility increases dramatically and the number of quarters that economic agents look back, to form their expectations, increases to ten quarters, as is suggested by the memory index. The message for central banks is that monetary policy has to be either more patient or more drastic. Importantly, we see from the memory index that the pattern of expectation formation has recently gone back to its pre-crisis level. From this analysis we may conclude that the local level model with SV (1)-(2) provides a flexible modelling framework because it takes into account stochastic changes in volatility, mean and persistence simultaneously. It is a powerful device to disentangle different volatility responses to macroeconomic shocks.

5.2 An extended analysis for monthly core inflation

To provide a more complete and detailed analysis for the dynamic structures in inflation, we consider the first difference of logarithm of the monthly US core Consumer Price Index (CPI; Bureau of Labor Statistics, with the city average CPI index; all items less food and energy; not seasonally-adjusted) between 1957:1 and 2015:1, we have $n = 695$. We take this time series as the indicator of monthly core inflation. The data are presented in Figure 4 from which we learn that core inflation is a noisy time series and is subject to trend and seasonal effects. Hence we extend the local level model with a seasonal component and obtain the model as described in Section 2.4. We have the decomposition model $y_t = \mu_t + \gamma_t + \epsilon_t$ with trend μ_t , seasonal γ_t and irregular ϵ_t , where each component is driven by independent sequences of disturbances which have mean zero and stochastically time-varying log-variances h_t^μ , h_t^γ and h_t^ϵ , respectively. The disturbance sequences are independent of each other but their stochastic volatility processes can be correlated with each other. Further details of the modelling framework are discussed in Section 2 and the simulation-based estimation methodology is presented in Section 3.




Figure 4

(i) US Core Monthly CPI and (ii) the first difference of log CPI (Core Inflation).

The estimation results for the UC and UCSV models are presented in Table 3. The estimated scaling coefficients for the UC models imply signal-to-noise ratios that are equal to the typical values of 0.089 for the trend and 0.035 for the seasonal. The overall fit is also reasonable. The reported p -values for the diagnostic test statistics are based on the standardised residuals and indicate strong rejection of the underlying null hypotheses. The top row of Figure 5 presents standardised residuals, correlograms of residuals and residuals squared, and scaled cumulative sum of squared residuals (CUSUM) for the UC model. These

Table 3
UCSV Estimation Results for Monthly Core Inflation

Parameter	UC	UC-D	UCSV	UCSV-D
α_y			-4.377 [-5.150, -3.604]	-4.346 [-5.057, -3.635]
ϕ_y			0.984 [0.975, 0.993]	0.975 [0.967, 0.983]
$\rho_{y\mu}$			0.638 [0.496, 0.780]	0.594 [0.426, 0.761]
σ_y	0.144 [0.133, 0.156]	0.131 [0.120, 0.142]	0.172 [0.126, 0.218]	0.182 [0.135, 0.228]
σ_μ	0.043 [0.031, 0.054]	0.043 [0.032, 0.053]	0.150 [0.082, 0.217]	0.149 [0.082, 0.216]
σ_γ	0.027 [0.020, 0.033]	0.029 [0.022, 0.037]	0.122 [0.092, 0.152]	0.109 [0.079, 0.139]
D_1 (1974:2)		0.419 [0.107, 0.731]		0.436 [0.123, 0.7450]
D_2 (1974:11)		-0.267 [-0.576, 0.042]		-0.282 [-0.587, -0.024]
D_3 (1980:7)		-1.213 [-1.521, -0.906]		-1.160 [-1.469, -0.852]
D_4 (1981:9)		-0.425 [-0.732, -0.117]		-0.413 [-0.723, -0.103]
D_5 (1982:8)		-0.421 [-0.730, -0.113]		-0.437 [-0.745, -0.128]
Normality	0.000	0.000	0.850	0.296
Box-Ljung	0.000	0.000	0.091	0.024
H(n/3)	0.000	0.000	0.339	0.137
LL	179.934	220.300	373.762	394.074

UC: trend, seasonal plus irregular model. UC-D: UC with dummies. UCSV: UC with SV in all components. UCSV-D: UCSV with dummies. The dummy variables account for some clear outliers in the inflation series. Scale coefficients σ_y , σ_μ and σ_γ are for SV processes h_t^y , h_t^μ and h_t^γ , respectively. For UC and UC-D, they denote the constant standard deviations of the trend, seasonal and irregular components, respectively. The values between square brackets indicate 95% confidence intervals. For the three tests, we present their p -values. LL is the maximised log-likelihood value.

diagnostic graphs clearly provide empirical evidence of volatility clustering and autocorrelation in the residuals. The breaks in the CUSUM plot show that the UC model fails to account for the economic shocks around 1980. The UC model appears to fail in capturing the key dynamic features in monthly US core inflation. Consequently, it may deliver unreliable forecasts and it cannot be used for detrending or seasonal adjustment. The column labeled UC-D in Table 3 shows that inclusion of dummy variables which account for events such as two oil crises and monetary policy change significantly improves the log-likelihood value. The test statistics, however, still indicate model misspecification.

The introduction of SV in our UC model increases the maximised log-likelihood even more when compared to the inclusion of the set of dummy variables. The estimated correlation between the innovations of h_t^y and h_t^μ has a positive sign and is larger in magnitude than the one estimated for quarterly inflation reported in the previous section. We specifically draw the attention on the diagnostic tests which are all satisfactory for the UCSV model, although the Box-Ljung Q test is marginally rejected at the 10% level. The diagnostic graphs for the standardised residuals of the UCSV model are presented in the bottom row of Figure 5. The




Figure 5
UCSV Residual Diagnostics:

- (i): UC standardised residuals; (ii) UC correlograms of residuals (solid lines) and squared residuals; (iii) UC cumulative sum of squared residuals;
(iv)-(vi) UCSV counterparts.

residuals resemble a white noise process. The serial correlation for the UCSV residuals is overall small; this is also found for the squared residuals. The CUSUM plot for the UCSV residuals further shows a significant improvement when compared to the CUSUM plot of the UC residuals. We may conclude from the estimation results and the residual diagnostics for the UCSV model that the SV component is an important feature of the monthly US core inflation. The UCSV model tracks time-varying unobserved components and their time-varying volatility and it therefore captures different economic shocks in a structural way and help decipher the response of inflation on these shocks. In the same way as for the UC model, we consider the robustness check of including a set of dummy variables. In case of the UCSV model, the maximised log-likelihood value increases only marginally while the parameter estimates do not change by much. Furthermore, despite the significant estimates for the outlier dummy coefficients, the test statistics do not suggest any further superiority towards the UCSV model without the set of dummy variables.

Figure 6 presents the smoothed estimates (signal extraction) of the components irregular, trend and seasonal, together with the corresponding smoothed estimates of the SV compo-

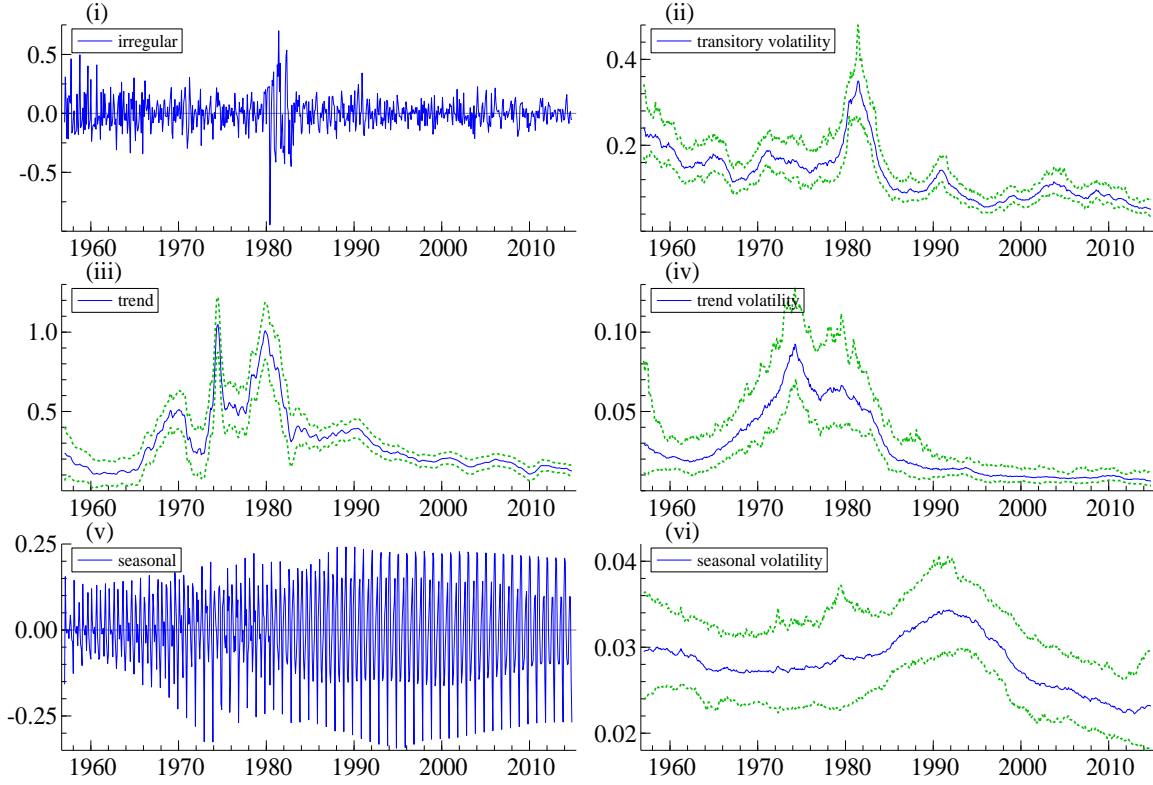


Figure 6

UCSV Decomposition of Monthly Core Inflation.

- (i) Irregular component ϵ_t ; (ii) Irregular volatility $\exp(h_t^y/2)$; (iii) Trend component μ_t ;
(iv) Trend volatility $\exp(h_t^\mu/2)$; (v) Seasonal component γ_t ; (vi) Seasonal volatility $\exp(h_t^\gamma/2)$.

nents, for the UCSV model. Our main findings are as follows. Firstly, as suggested by the asymptotic 95% confidence bands, the seasonal stochastic volatility is estimated almost as a constant variance over time, with a mild peak in the early 1990s and a decline afterwards. Secondly, the estimated irregular (transitory) volatility, that is $E\left(\exp(\frac{1}{2}h_t^y)|y_{1:n}; \psi = \hat{\psi}\right)$, shows a clear increase in the early 1980s when the monetary regime change has started. However, the irregular volatility signal has been slowly declining overall in the last sixty years. Thirdly, we detect a clear structural break in the trend volatility which we associate with the “Great Moderation”. Its peak is observed approximately in 1974 when the first oil crisis took place, after which a distinctively wider confidence bands of the trend volatility is observed around 1979 when the second oil crisis happened. We find that the UCSV model is effective in capturing the different and distinguished volatility regimes in the three dynamic components. Also these findings corroborate those from the previous section on quarterly

headline inflation: the major contributor of inflation volatility response to monetary policy change is the transitory volatility, while volatility response to energy shocks comes from trend volatility. Such findings are also discussed in other empirical studies; see, for example, [Goodfriend and King \(1997\)](#), [Erceg et al. \(2000\)](#) and [Aoki \(2001\)](#).

5.3 Forecast analysis

To investigate the forecasting performances of the UC and UCSV models, we carry out a real-time out-of-sample forecasting exercise, with both point and density forecasts. We consider the two main models and two variations of the UCSV models (a model with only the transitory SV and a model with only SV in the trend component).

The forecasting study is designed as follows. We consider a first sub-sample of the monthly inflation time series, we take the period from 1957:1 to 1985:1, and we consider the four models and estimate their parameters using the time series observations from this sub-sample. Based on the estimated parameters, we compute the point forecasts for each model, h -steps ahead for $h = 1, 3, 6, 12$ months, using the particle filter. We compare the forecasts with the corresponding actual (future) observations and compute the forecast error. We repeat these steps for the next sub-sample of the time series, from 1957:1 to 1985:2. We repeat these steps until the end of the sample (more precisely, until the forecasted observation is still in the sample). This is the design of our rolling window out-of-sample forecasting exercise. When all h forecast errors are recorded for all models and all rolling windows, we compute three summary forecast precision statistics: mean forecast error (MFE_h), mean absolute forecast error (MAFE_h) and root mean squared error (RMSE_h) which are given by

$$\text{MFE}_h = \frac{1}{k} \sum_t v_{t,h}, \quad \text{MAFE}_h = \frac{1}{k} \sum_t |v_{t,h}|, \quad \text{RMFE}_h = \sqrt{\frac{1}{k} \sum_t v_{t,h}^2},$$

where $v_{t,h} = y_{t+h} - \hat{y}_{t+h|t}$ is the h -step forecast error at time t with $\hat{y}_{t+h|t} = E(y_{t+h}|y_{1:t})$, and where k is the number of available forecasts. The results for our time series of monthly US core inflation are presented in [Table 4](#).

The results in [Table 4](#) strongly suggest that the forecast accuracy achieved by the UCSV model is much higher when SV components are included in the model. When considering the

Table 4
UCSV Point Forecast Evaluations

	MFE _h				MAFE _h				RMSE _h			
	<i>h</i> = 1	<i>h</i> = 3	<i>h</i> = 6	<i>h</i> = 12	<i>h</i> = 1	<i>h</i> = 3	<i>h</i> = 6	<i>h</i> = 12	<i>h</i> = 1	<i>h</i> = 3	<i>h</i> = 6	<i>h</i> = 12
with SV	<i>h</i> = 1	<i>h</i> = 3	<i>h</i> = 6	<i>h</i> = 12	<i>h</i> = 1	<i>h</i> = 3	<i>h</i> = 6	<i>h</i> = 12	<i>h</i> = 1	<i>h</i> = 3	<i>h</i> = 6	<i>h</i> = 12
–	-0.006	-0.005	-0.009	-0.014	0.109	0.117	0.108	0.098	0.144	0.153	0.146	0.141
h_t^y	-0.004	-0.004	-0.010	-0.011	0.091	0.109	0.098	0.092	0.137	0.142	0.140	0.137
h_t^μ, h_t^γ	-0.005	-0.006	-0.008	-0.009	0.092	0.113	0.096	0.094	0.134	0.150	0.141	0.137
$h_t^y, h_t^\mu, h_t^\gamma$	-0.004	-0.004	-0.008	-0.010	0.087	0.097	0.094	0.092	0.128	0.131	0.124	0.122

Comparisons of point forecast performances for four different UCSV model specifications: without SV, the UC model (–); with only transitory SV (h_t^y); without transitory SV (h_t^μ, h_t^γ); and with all SV. We present the results for monthly ($h = 1$), quarterly ($h = 3$), semiannual ($h = 6$), and annual forecasts ($h = 12$).

models with none or a smaller set of SV components, they show a more mixed performance for the different forecast horizons. For the forecast horizons $h = 6$ and 12, we obtain MAFE and RMSE statistics that are lower than for the shorter horizons. It shows that the point forecasts show a relatively strong seasonal effect. Our overall conclusion is that the forecast of the UCSV model with all SV components included is superior in its forecasting performance for all considered forecast horizons.

Based on the same rolling window out-of-sample forecasts, we consider the DM statistic of Diebold and Mariano (2012) to verify whether the density forecasts delivered by the full UCSV model is the best among the other specifications. For this purpose, we define r_t^l and r_t^u as the lower and upper bounds of the tolerance band, respectively; we can take r_t^l and r_t^u as the empirical quantiles of the inflation series of the t^{th} rolling window with its empirical distribution function denoted by \tilde{F}_t . It is then implied that $\tilde{F}_t(r_t^u) - \tilde{F}_t(r_t^l) = I$ where I is the probability mass within the tolerance band; in our study we consider the three values $I = 0.4, 0.6, 0.8$. For the actual verification of our density forecasts, we consider three proper scoring rules where *proper* refers to the assumption that the scoring rule always yields the highest score for the unknown underlying data generation process of the time series. We adopt three scoring rules: weighted probability score (S^{wps}), conditional likelihood score (S^{cl}), and censored likelihood score (S^{csl}); these are formally defined in Diks, Panchenko, and Van Dijk (2011) and shown in the supplementary appendix.

Table 5 compares the density forecast performance of the full UCSV with the other model specifications, conditional on the chosen region of interest as indicated by $I = 0.8, 0.6, 0.4$. Many positive and significant score values are reported and they imply that the null of equal

Table 5
UCSV Density Forecast Evaluations

with SV	$I = 0.8$				$I = 0.6$				$I = 0.4$			
S^{wps}	$h = 1$	$h = 3$	$h = 6$	$h = 12$	$h = 1$	$h = 3$	$h = 6$	$h = 12$	$h = 1$	$h = 3$	$h = 6$	$h = 12$
–	1.765 [0.078]	3.476 [0.001]	2.540 [0.011]	2.079 [0.038]	1.431 [0.152]	2.640 [0.008]	1.922 [0.055]	1.644 [0.100]	2.736 [0.006]	3.511 [0.000]	1.625 [0.001]	0.086 [0.101]
h_t^y	3.303 [0.001]	-0.640 [0.522]	-0.025 [0.980]	0.039 [0.969]	-0.790 [0.429]	1.042 [0.297]	-0.015 [0.988]	0.018 [0.985]	-0.196 [0.844]	-0.798 [0.425]	-0.025 [0.980]	0.029 [0.977]
h_t^μ, h_t^γ	2.355 [0.019]	3.214 [0.001]	1.138 [0.255]	0.600 [0.549]	1.358 [0.174]	1.051 [0.293]	0.651 [0.515]	0.322 [0.748]	3.109 [0.002]	2.189 [0.029]	0.641 [0.521]	0.428 [0.669]
S^{cl}												
–	1.143 [0.253]	1.890 [0.059]	0.939 [0.348]	0.912 [0.362]	0.142 [0.887]	2.103 [0.036]	0.515 [0.607]	0.410 [0.681]	1.773 [0.076]	1.986 [0.047]	0.837 [0.402]	0.573 [0.567]
h_t^y	2.413 [0.016]	-0.759 [0.447]	0.423 [0.672]	-0.196 [0.848]	1.228 [0.220]	0.228 [0.820]	0.199 [0.843]	-0.105 [0.917]	0.048 [0.961]	1.625 [0.104]	0.279 [0.780]	-0.098 [0.921]
h_t^μ, h_t^γ	1.880 [0.060]	1.906 [0.057]	0.737 [0.460]	0.609 [0.543]	1.119 [0.263]	1.285 [0.199]	0.509 [0.611]	0.279 [0.781]	1.999 [0.046]	1.224 [0.221]	0.443 [0.658]	0.378 [0.706]
S^{csl}												
–	3.044 [0.002]	2.633 [0.009]	1.920 [0.055]	1.216 [0.224]	4.168 [0.000]	3.128 [0.002]	1.149 [0.251]	0.657 [0.511]	2.386 [0.017]	2.791 [0.005]	1.089 [0.276]	0.722 [0.470]
h_t^y	2.180 [0.029]	1.803 [0.071]	0.292 [0.770]	-0.525 [0.599]	1.814 [0.070]	-0.187 [0.852]	0.134 [0.894]	-0.378 [0.705]	0.599 [0.549]	-0.534 [0.593]	0.151 [0.880]	-0.479 [0.632]
h_t^μ, h_t^γ	1.291 [0.197]	2.365 [0.018]	1.051 [0.293]	1.348 [0.178]	2.316 [0.021]	1.929 [0.054]	0.651 [0.515]	0.699 [0.485]	1.940 [0.052]	2.577 [0.010]	0.935 [0.350]	0.940 [0.347]

The likelihood-based scoring tests S^{wps} , S^{cl} and S^{csl} indicate the significance of the difference between the forecast precision of the full UCSV model and three other model specifications: without SV, the UC model (–); with only transitory SV (h_t^y); without transitory SV (h_t^μ, h_t^γ). We present the results for monthly ($h = 1$), quarterly ($h = 3$), semiannual ($h = 6$), and annual forecasts ($h = 12$). The aim is to verify whether the full UCSV model significantly outperforms the three other model specifications in terms of density forecast; the p -values are in the square brackets. A significant positive value of the scoring test indicates the superiority of the full UCSV model when compared to the other UCSV specification. A negative value indicates the superiority of the specific UCSV model against the full model.

predictive ability between the full UCSV and another model specification is strongly rejected in favour of the former model. The short forecast horizons $h = 1, 2$ show the superiority of the full UCSV model convincingly for almost all scoring tests S and all regions I . The UCSV model with only h_t^y is the most competitive as some scoring tests are not significant, especially for $I = 0.6, 0.4$. For the longer forecast horizons, the evidence is more mixed but overall the UCSV model outperforms the UC model convincingly. In particular, all three score test values are positive for all horizons and a large selection of these are significant. Finally, for $I = 0.8$ and 0.6 , the full UCSV model produces better density forecast than the UCSV model with only h_t^y for all three scoring rules. We may conclude that the volatility h_t^μ (and h_t^γ) should be modelled on top of transitory volatility. But when we narrow the region of interest to $I = 0.4$, the full and only h_t^y UCSV models deliver statistically equally accurate density forecasts. The overall conclusion is therefore that the full UCSV model shows overall superior forecasting performance for monthly US core inflation.

6 Conclusion

We have considered the important challenge of modelling inflation volatility responses to energy and monetary shocks which is of key importance for monetary policy analysis that includes inflation targeting and the construction of inflation bands. Such monetary policy tools are typically based on inflation forecasts. We have generalised the local level model with stochastic volatility of [Stock and Watson \(2007, 2008\)](#) by including dynamic seasonal and cycle components with stochastic volatility. We further have proposed a likelihood-based estimation procedure using the importance sampling method. It is a fast estimation method and therefore a competitive alternative to the existing Bayesian methods. In a Monte Carlo study we have shown the accuracy of our proposed parameter estimation method; its asymptotic validity is shown by the verification of the finite variance assumption of [Geweke \(1989\)](#). In our empirical study, we have confirmed earlier findings that quarterly US inflation has experienced an increasing persistence during the financial crisis but that it has returned to its pre-crisis level in the recent years. Furthermore, based on monthly US core inflation, we have found that stochastic volatility is an important feature in inflation. The transitory volatility appears to respond to the Volcker-Greenspan monetary regime, while the trend volatility is more responsive to the two oil crises in the 1970s. Finally, point and density forecasting exercise provides evidence that our proposed model is able to produce reliable inflation forecast at different horizons.

References

- Aoki, K. (2001). Optimal monetary policy responses to relative-price changes. *Journal of Monetary Economics* 48(1), 55–80.
- Atkeson, A. and L. E. Ohanian (2001). Are Phillips curves useful for forecasting inflation? *Federal Reserve bank of Minneapolis quarterly review* 25(1), 2–11.
- Bank of England (2013, August). Inflation report. *Monetary policy trade-offs and forward guidance*.

- Cecchetti, S. G. and G. Debelle (2006). Has the inflation process changed? *Economic Policy* 21(46), 312–352.
- Cecchetti, S. G., P. Hooper, A. K. Kashyap, and K. L. Schoenholtz (2017, March). Deflating inflation expectations: The implications of inflation’s simple dynamics. In *U.S. Monetary Policy Forum*.
- Chen, R. and J. S. Liu (2000). Mixture Kalman filters. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 62(3), 493–508.
- Clarida, R., J. Gali, and M. Gertler (2000). Monetary policy rules and macroeconomic stability: Evidence and some theory. *Quarterly Journal of Economics*, 147–180.
- Congressional Budget Office (2016, January). Budget projections. *The Budget and Economic Outlook: 2016 to 2026*.
- Constâncio, V. (2016). The challenge of low real interest rates for monetary policy. In *Macroeconomics Symposium at Utrecht School of Economics*, Volume 15.
- Cúrdia, V. et al. (2015). Why so slow? A gradual return for interest rates. *FRBSF Economic Letter* 12(2015), 43–45.
- Danielsson, J. (1994). Stochastic volatility in asset prices estimation with simulated maximum likelihood. *Journal of Econometrics* 64(1), 375–400.
- Del Negro, M., D. Giannone, M. P. Giannoni, and A. Tambalotti (2017, May). Safety, liquidity, and the natural rate of interest. *FRBNY Staff Reports* 812.
- Diebold, F. X. and R. S. Mariano (2012). Comparing predictive accuracy. *Journal of Business & Economic Statistics*.
- Diks, C., V. Panchenko, and D. Van Dijk (2011). Likelihood-based scoring rules for comparing density forecasts in tails. *Journal of Econometrics* 163(2), 215–230.
- Durbin, J. and S. J. Koopman (1997). Monte Carlo maximum likelihood estimation for non-Gaussian state space models. *Biometrika* 84(3), 669–684.

- Durbin, J. and S. J. Koopman (2012). *Time series analysis by state space methods*. Number 38. Oxford University Press.
- Durham, G. B. and A. R. Gallant (2002). Numerical techniques for maximum likelihood estimation of continuous-time diffusion processes. *Journal of Business & Economic Statistics* 20(3), 297–338.
- Erceg, C. J., D. W. Henderson, and A. T. Levin (2000). Optimal monetary policy with staggered wage and price contracts. *Journal of Monetary Economics* 46(2), 281–313.
- Federal Open Market Committee (2009). Greenbook. *Current Economic and Financial Conditions: Summary and Outlook*.
- Geweke, J. (1989). Bayesian inference in econometric models using Monte Carlo integration. *Econometrica: Journal of the Econometric Society*, 1317–1339.
- Goodfriend, M. and R. King (1997). The new neoclassical synthesis and the role of monetary policy. In *NBER Macroeconomics Annual 1997, Volume 12*, pp. 231–296. MIT Press.
- Gordon, R. J. (1990). The Phillips curve now and then. Technical report, National Bureau of Economic Research, Inc.
- Harvey, A. C. (1989). *Forecasting, structural time series models and the Kalman filter*. Cambridge university press.
- Holston, K., T. Laubach, and J. C. Williams (2017). Measuring the natural rate of interest: International trends and determinants. *Journal of International Economics*.
- Jäckel, P. (2005). A note on multivariate Gauss-Hermite quadrature.
- King, R. and A. L. Wolman (1999). What should the monetary authority do when prices are sticky? In *Monetary Policy Rules*, pp. 349–404. University of Chicago Press.
- Koopman, S. J., A. Lucas, and M. Scharth (2015). Numerically accelerated importance sampling for nonlinear non-Gaussian state-space models. *Journal of Business & Economic Statistics* 33(1), 114–127.

- Koopman, S. J., N. Shephard, and D. Creal (2009). Testing the assumptions behind importance sampling. *Journal of Econometrics* 149(1), 2–11.
- Laubach, T. and J. C. Williams (2003). Measuring the natural rate of interest. *The Review of Economics and Statistics* 85(4), 1063–1070.
- Lawrence, J. C. and G. D. Rudebusch (2016). A new normal for interest rate? Evidence from inflation-indexed debt. *Mimeo, FRBSF Working Paper*.
- Marques, C. R. et al. (2004). Inflation persistence: facts or artefacts? Technical report, European Central Bank.
- Pivetta, F. and R. Reis (2007). The persistence of inflation in the United States. *Journal of Economic dynamics and control* 31(4), 1326–1358.
- Platanioti, K., E. McCoy, and D. Stephens (2005). A review of stochastic volatility: univariate and multivariate models. Technical report, working paper, Department of Mathematics, Imperial College London.
- Richard, J.-F. and W. Zhang (2007). Efficient high-dimensional importance sampling. *Journal of Econometrics* 141(2), 1385–1411.
- Shephard, N. (2005). Stochastic volatility. *Economics Group, Nuffield College, University of Oxford, Economics Papers*.
- Shephard, N. (2015). Martingale unobserved component models. In S. J. Koopman and N. Shephard (Eds.), *Unobserved Components and Time Series Econometrics*, pp. Chapter 10. Oxford University Press.
- Shephard, N. and M. K. Pitt (1997). Likelihood analysis of non-Gaussian measurement time series. *Biometrika* 84(3), 653–667.
- Stock, J. H. and M. W. Watson (2007). Why has US inflation become harder to forecast? *Journal of Money, Credit and banking* 39(s1), 3–33.
- Stock, J. H. and M. W. Watson (2008). Phillips curve inflation forecasts. Technical report, National Bureau of Economic Research.